

Claims

1. Method of evaluating a query involving one or more relational databases, each comprising a relational database management system (RDBMS), said query relating to at least two tables of said relational database,
said method comprising
determining a set of tables, including the tables referred to in the query, wherein each of the tables in said set is linked to at least one other table, such that, in a graphical representation of the database wherein the tables are represented as nodes and links between the tables are represented as lines between the nodes, they form a connected graph connecting nodes corresponding to the tables referred to in the query,
performing a query on said set equivalent to the initial query, wherein said step of performing said query on said set comprises performing consecutive partial queries, wherein a result of a previous query is used as input for a later query,
combining the results of said partial queries to obtain a result to the initial query.
2. Method according to claim 1, wherein said step of performing a query on said set comprises:
determining two or more partial queries on said set of tables which, taken in conjunction, are equivalent to the initial query, wherein there is at least a first partial query involving a first subset of said set of tables and a second partial query involving a second subset of tables at least partly distinct from said first subset,
wherein at least one table in said first subset not contained in said second subset has a link to a table in said second subset and not contained in said first subset, said link corresponding to a line in said graph the removal of which would render the graph unconnected,
wherein said first query renders a result comprising values of the link key between said two tables and wherein said second query has as input the values of the link key determined in the first query,
performing said first query,
performing said second query, using the result of said first query,
combining the result of all partial queries to render a result to the initial query.

3. Method according to claim 2, wherein said first and second subset are distinct from each other.
4. Method according to one of claims 1 to 3, wherein said first and second partial queries are dynamically created during runtime.
5. Method according to one of claims 1 to 4, wherein said result of a previous query comprises the value of a foreign key of a table involved in a later query and wherein said value of said foreign key is used as input for said later query.
6. Method according to one of claims 1 to 5, wherein the result of said partial queries is stored as an object or objects.
7. Method according to one of claims 1 to 6, wherein after each partial query a redundancy check and/or a check for consistency is carried out in the respective result and a result of said partial query purged of redundancies and/or inconsistencies is stored.
8. Method according to claims 6 and 7, wherein said redundancy and/or check for consistency is carried out in creating said object comprising the result of said partial query or directly on said object after the creation of the same.
9. Method according to one of claims 1 to 8, wherein each partial query involves a table or a plurality of tables linked to each other and wherein each partial query other than a first has as input previously established values of link keys, said link keys linking said table or one or more of said plurality of tables to another table not involved in said partial query.
10. Method according to one of claims 1 to 9, wherein said graph comprises at least one branch node having links to at least two other nodes and wherein tables referred to in the initial query are related to separate branches deriving from said branch node, wherein a partial query is carried out involving the table corresponding to said branch node (branch table) and wherein at least one partial query is carried out for one or

more tables contained in each branch which has the result of the partial query involving the branch table as an input.

11. Method according to one of claims 1 to 10, wherein said tables contained in said set are not comprised in one single database.
12. Method of evaluating a query involving at least one relational database comprising a relational database management system (RDBMS), said query relating to at least one table of said relational database,
said method comprising
determining a table of said relational database as a gateway table for evaluating said query,
retrieving one or more unique identifiers of said gateway table related to one or more entries in a table to be queried,
retrieving information from one or more tables to be queried related to said retrieved unique identifiers of said gateway table,
providing a result to said query,
wherein at least one of said steps of retrieving one or more primary keys of said gateway table or said step of retrieving information from tables to be queried related to said retrieved primary keys of said gateway table comprises the evaluation of a query according to one of claims 1 to 8.
13. Method according to claim 9, wherein a first of said partial queries involves a gateway table and a query other than the first query relates to a table referred to in the initial query.
14. Method according to claim 12 or 13, wherein said relational database comprises one or more predetermined hub tables and said query relates to at least one table of said relational database and wherein said method comprises:
retrieving one or more unique identifiers of a hub table related to one or more entries in a table to be queried,
retrieving information from tables to be queried related to said retrieved unique identifiers of said hub table,
providing a result set to said query.

15. Method according to one of claims 12 to 14, wherein at least one library is defined on one or more of said databases, said library consisting of tables linked to each other and having exactly one table defined as hub table.
16. Method according to claims 12 to 15, wherein said query is for complete sets of related entries of said relational database or of a library or for parts of such complete sets of entries and comprises one or more query conditions related to said database, wherein said method comprises:
 - identifying a gateway table related to entries specified in a query condition,
 - identifying one or more unique identifiers of said gateway table related to said entries conforming to query conditions,
 - retrieving a complete set of related entries or part thereof related to said primary keys of said gateway tables.
17. Method according to one of claims 14 to 16, wherein said query involves at least a second searchable entity outside said database or outside a library involved in said query, said second entity comprising sub-entities each having at least one identifier uniquely identifying said sub-entities, and wherein said method comprises:
 - retrieving one or more identifiers of sub-entities of said second searchable entity related to said query,
 - retrieving one or more unique identifiers of a hub table of said relational database or library related to said retrieved identifiers of said sub-entities,
 - retrieving sets of related entries or predetermined parts thereof related to said retrieved unique identifiers of said hub,
 - retrieving information from a sub-entity identified by a retrieved identifier in said second entity,
 - combining the retrieved information from said second searchable entity and said database or library into a result.
18. Method according to one of claims 14 to 17, wherein said query involves at least a second searchable entity outside said database or outside a library involved in said query and comprising sub-entities, each sub-entity having at least one identifier uniquely specifying said sub-entity, and wherein said method comprises:

- retrieving one or more unique identifiers of a hub table of said database or library related to entries related to said query,
 - retrieving identifiers of sub-entities of said second searchable entity related to said retrieved unique identifiers of said hub table;
 - retrieving sets of related entries or predetermined parts thereof related to said retrieved unique identifiers of said hub,
 - retrieving information from said sub-entities identified by identifiers retrieved in said second searchable entity,
 - combining the retrieved information from said second entity and said database or library into a result.
19. Method according to one of claims 17 or 18, wherein said second searchable entity is a second relational database or a library and said identifier is a unique identifier of a hub table in said second relational database or library.
20. Method according to one of claims 17 or 18, wherein said second searchable entity is a collection of flat files with the sub-entities being flat files in this collection.
21. Method according to one of claims 17 to 20, wherein said step of retrieving a relation between identifiers of said second searchable entity and unique identifiers of hubs of said database or library comprises the step of discarding combinations of identifiers of hubs with identifiers of said second searchable entity which are not consistent with the query conditions and retrieving only such additional information related to an identifier which is comprised in a combination of identifiers consistent with the selection parameters.
22. Method according to one of claims 12 to 21, wherein the query relates to tables related to at least two hub tables, wherein said method comprises:
- retrieving one or unique identifiers of a hub table or hub tables, said identifier being related to entries satisfying query conditions in tables related to the respective hub,
 - retrieving unique identifiers of the respective other hub or hubs related to said retrieved unique identifiers related to entries satisfying the query conditions,

- retrieving sets of related entries or predetermined parts thereof which are related to said retrieved unique identifiers of said hubs according to the query,
 - combining the retrieved information related to said hubs into a result.
23. Method according to claim 22, wherein said step of retrieving a relation between unique identifiers of said hub tables comprises the step of discarding combinations of unique identifiers of hub tables which are not consistent with the query conditions and retrieving only such additional information related to a unique identifier of at least one hub table which is comprised in combinations of unique identifiers consistent with the search parameters.
24. Method according to claim 22 or 23, wherein at least one of said hubs is the hub of a library and the query relates to said library.
25. Method according to claim 24, wherein the query relates to two libraries and said hubs are hubs of two libraries.
26. Method according to one of claims 22 to 25, wherein said two hubs are hubs within the same relational database.
27. Method according to one of claims 12 to 26, wherein the step of retrieving unique identifiers of a hub table and/or identifiers of a searchable entity which are related to another unique identifier of a hub table and/or identifier of a searchable entity is performed on the basis of pre-established relations between keys and/or identifiers of said entity.
28. Method according to one of claims 12 to 26, wherein the step of retrieving unique identifiers of a hub table and/or identifiers of a searchable entity which are related to another unique identifier of a hub table and/or identifier of a searchable entity is performed dynamically during the execution of the query.
29. Method according to one of claims 12 to 28, wherein in performing said step of retrieving information related to a unique identifier of said gateway table, selected tables are queried which, in a graphical representation of the database wherein the tables are

represented as nodes and links between the tables are represented as lines between the nodes, form a connected graph connecting the gateway table to tables referred to in the initial query.

30. Method according to claim 29, wherein said step of querying tables on said graph comprises performing consecutive partial queries, wherein a result of a previous query is used as input for a later query, a first of said partial queries involving the gateway table and a query other than the first query relating to a table referred to in the initial query.
31. Method according to one of claims 12 to 30 comprising the steps of
 - identifying the hub table or hub tables related to tables referred to in the initial query,
 - determining in said graphical representation of said database, at least for one hub table, an optimum graph connecting said hub to all tables which are related to said hub and referred to in the initial query,
 - performing consecutive partial queries involving tables which are consecutive to each other with regard to said optimum graph.
32. Method according to one of claims 12 to 31, wherein said step of retrieving primary keys of said gateway table comprises:
 - determining a table that is referred to in the initial query,
 - determining, in a graphical representation of said database, wherein tables are represented as nodes and links between tables as lines between the nodes, a gateway table connected to said table,
 - querying said database for one or more indices of the gateway table which are related to said table.
33. Method according to claim 32, wherein one or more specific entries of said table are implied by a query condition and said database is queried for indices of said gateway table which are related to said entry or entries.
34. Method according to claim 32 or 33, wherein in said graphical representation, a path from said table to said gateway table is established and said query for said indices is

performed by querying all tables corresponding to nodes in said graph for the values of link keys between the tables in said graph, starting from the table referred to in the query and, given the case, certain entries thereof.

35. Method according to claim 34, wherein said path is selected as a shortest path between said table and said gateway table according to a predetermined metric.
36. Method according to claim 34 or 35, wherein said path is part of or identical to the graph for determining partial queries for retrieving additional information from tables related to said gateway table.
37. Method according to one of claims 32 to 36, comprising the steps of:
 - determining a unique identifier for one or more rows of the gateway table related to said indices, if an index or a group of indices related to the same row of the gateway table determined by said step of querying the database does not uniquely identify a row of said gateway table.
38. Method according to one of claims 12 to 37, wherein partial queries used for evaluating the initial query are at least partially created dynamically during the process of said evaluation.
39. Method according to one of claims 1 to 38, wherein a result is represented in an object oriented representation.
40. Method according to claim 39, wherein the result of said initial query is expressed as an object derived by means of object-relational mapping.
41. Method according to one of claims 1 to 40, wherein said evaluation of said query is performed under the control of an object manager, said object manager comprising a sequence of commands to be executed by a computer system.
42. Method according to claim 41, wherein said object manager handles an object which represents the schema or part of a schema of one or more databases to be queried.

43. Method according to claim 41 or 42, wherein said object manager defines classes which are dynamically created and instantiated.
44. Data processing system for controlling the evaluation of a query involving a relational database comprising a relational database management system (RDBMS), said query relating to at least one table of said relational database, comprising:
means for determining a set of tables, including the tables referred to in the query, wherein each of the tables in said set is linked to at least one other table, such that, in a graphical representation of the database wherein the tables are represented as nodes and keys linking the tables are represented as lines between the nodes, they form a connected graph connecting nodes corresponding to the tables referred to in the query,
means for performing or causing to be performed a query on said set equivalent to the initial query, wherein performing said query on said set comprises performing consecutive partial queries, wherein a result of a previous query is used as input for a later query,
means for combining or causing to be combined the results of said partial queries to obtain a result to the initial query.
45. Data processing system according to claim 40, comprising means for controlling the execution of a method according to one of claims 1 to 43 by a data processing system or data processing systems.
46. Computer program causing a computer or computer system, when executed thereon, to perform the steps of a method according to one of claims 1 to 43.
47. A computer readable storage medium, comprising a program according to claim 46.